

Module Eight

This paper presents a MAC policy for an MLS relational DBMS with views as the security objects, discusses advantages and disadvantages of using views as the security objects, and describes an implementation approach.

- Wisema86 Wiseman, S., "A Secure Capability Computer System," *Proceedings of the 1986 IEEE Symposium on Security and Privacy*, pp. 86-94, April 1986.

A secure computer system based on a capability architecture is described. Abstract types are used to provide separation and the reference monitor function. By providing a trusted path from the user to security critical operations, full DAC and MAC is enforced.

- Woodwa87 Woodward, J.P.L., "Exploiting the Dual Nature of Sensitivity Labels," *Proceedings of the 1987 IEEE Symposium on Security and Privacy*, pp. 23-31, April 1987.

This paper describes the MAC scheme for the Compartmented Mode Workstation (CMW). The CMW assigns both a MAC level and a sensitivity label to each object in an attempt to keep objects from being under or over- classified.

Module Eight

- Meadow87 Meadows, C., "The Integrity Lock Architecture and its Application to Message Systems: Reducing Covert Channels," *Proceedings of the 1987 IEEE Symposium on Security and Privacy*, pp. 212-219, April 1987.
- This paper discusses the covert channel problem that can arise by leaking incorrect but unclassified data in order to covertly reveal classified data in an integrity-lock database management system (DBMS).
- Millen87 Millen, J.K., "Covert Channel Capacity," *Proceedings of the 1987 IEEE Symposium on Security and Privacy*, pp. 60-66, April 1987.
- "This paper establishes a connection between Shannon's theory of communication and information flow models that view a reference monitor as a state-transition automaton. The channel associated with a machine and a compromise policy is defined, and the capacity of that channel is taken as a measure of covert channel information rate."
- Parenty89 Parenty, T.J., "The Incorporation of Multi-Level IPC into UNIX," *Proceedings of the 1989 IEEE Symposium on Security and Privacy*, pp. 94-101, May 1989.
- Through the use of alternate forms of process synchronization and modification to UNIX semantics, it is possible to provide multi-level IPC in UNIX. These features facilitate the use of UNIX as a base for multi-level applications.
- Tsai87 Tsai, C., Gligor, V.D., and Chandrasekaran, C.S., "A Formal Method for the Identification of Covert Storage Channels in Source Code," *Proceedings of the 1987 IEEE Symposium on Security and Privacy*, pp. 74-87, April 1987.
- A formal method for the identification of covert storage channels is presented and its application to the source code of the Secure Xenix kernel is illustrated.
- Tsai88 Tsai, C. and Gligor, V.D., "A Bandwidth Computation Model for Covert Storage Channels and its Applications," *Proceedings of the 1988 IEEE Symposium on Security and Privacy*, pp. 108-121, April 1988.
- This paper describes a tool that determines the factors that affect the bandwidth of covert storage channels, which enables the placement of dynamically adjustable delays in multiprogrammed systems, which guarantees minimum performance impact with maximum reduction of covert channel bandwidth.
- Wilson88 Wilson, J., "Views as the Security Objects in a Multilevel Secure Relational Database Management System," *Proceedings of the 1988 IEEE Symposium on Security and Privacy*, pp. 70-84, April 1988.

Module Eight

- Young86 Young, W.D., Telega, P.A., and Boebert, W.E., "A Verified Labeler for the Secure Ada Target," *Proceedings of 9th National Computer Security Conference*, pp. 55-61, September 1986.

This paper describes the specification and verification of a prototype line printer labeler for the Secure Ada Target (SAT) machine. Used for labeling human readable output of the SAT.

Other Readings

- Bransta89 Branstad, M., Tajalli, H., Meyer, F., and Dalva, D., "Access Mediation in a Message Passing Kernel," *Proceedings of the 1989 IEEE Symposium on Security and Privacy*, pp. 66-72, May 1989.

This paper discusses how mandatory access mediation and discretionary access mediation are performed in the Trusted Mach kernel, a system that uses message passing as its primary means of communication both between tasks and within the kernel.

- Denning86 Denning, D.E., Akl, S.G., Morgenstern, M., Neumann, P.G., Schell, R.R., and Heckman, M., "Views for Multilevel Database Security," *Proceedings of the 1986 IEEE Symposium on Security and Privacy*, pp. 156-172, April 1986.

This paper describes basic view concepts for a multi-level security (MLS) relational database model. The model treats stored and derived data uniformly within the database schema. All data in the database is classified according to views called classification constraints which specify security levels for related data.

- Graubar88 Graubart, R.D., "Dual Labels Revisited," *Proceedings of the 4th Aerospace Computer Security Applications Conference*, pp. 167-172, December 1988.

This paper discusses the utility and need for trusted labels employed for functions other than access control. It describes the rationale for adopting two trusted labels in intelligence community security requirements. It advocates that dual labels are viable and necessary.

- Haigh86 Haigh, J.T., Kemmerer, R.A., McHugh, J., and Young, W.D., "An Experience Using Two Covert Channel Analysis Techniques On a Real System Design," *Proceedings of the 1986 IEEE Symposium on Security and Privacy*, pp. 14-24, April 1986.

This paper examines the application of two covert channel analysis techniques to a high level design for a real system - the Honeywell Secure Ada Target (SAT). The techniques used were a version of the non-interference model of MLS security due to Goguen and Meseguer and the shared resource matrix method of Kemmerer. Both techniques were applied to the Gypsy abstract model of the SAT.

Module Eight

National Computer Security Conference, pp. 47-54, September 1986.

This paper describes the functions of SE/VMS that support user registration and login, device and volume management, file creation and access, and the production of labeled printed output. The techniques that were used to implement SE/VMS are discussed.

- COVERT93 National Computer Security Center, *A Guide to Understanding Covert Channel Analysis of Trusted Systems*, NCSC-TG-030, Version 1, November 1993.

This document presents the relative merits of covert channel identification methods and of covert channel information sources, recommends sound bandwidth determination and handling policies and methods based on the TCSEC requirements, and defines the types of evidence that should be provided for handling assurance.

- Loepere84 Loepere, K., *Resolving Covert Channels within a B2 Class Secure System*, Multics Development Center -- Honeywell Information Systems, 1984.

This paper examines the problem of covert channels on Multics and attempts to analyze and resolve them relative to satisfying the B2 security requirements.

- Rubin90 Rubin, C., "UNIX System V with B2 Security," *Proceedings of the 13th National Computer Security Conference*, pp. 1-9, October 1990.

This paper describes the feature changes needed for UNIX System V to meet the TCSEC B2 requirements. Pages 4 and 5 specifically discuss the MAC mechanism and the new multilevel directories used to meet the MAC requirements.

Supplemental Readings

- Kemmer83 Kemmerer, R., "Shared Resource Matrix Methodology: An Approach to Identify Storage and Timing Channels," *ACM Transactions on Computer Systems*, Vol. I, No. 3, pp. 256-277, August 1983.

This paper outlines a practical methodology for discovering storage and timing channels that can be used through all phases of the software life cycle to increase the assurance that all channels have been identified. The methodology is presented and its application to three different descriptions (English, formal specification, and high order language) are discussed. This paper can also be found in a similar form in the proceedings of the 1982 IEEE conference on Security and Privacy.

Module Eight

Required Readings

TCSEC85 National Computer Security Center, *Department of Defense Trusted Computer System Evaluation Criteria*, DoD 5200.28-STD, December 1985.

Sections 3.1.1.4, 3.2.1.4, 3.3.1.4, and 4.1.1.4 contain the MAC requirements, which are summarized on pages 102-103. Additional information about MAC is supplied in Sections 5.3.1.1, 5.3.1.3, 7.3.1, 7.3.2, 8.0, and 9.0 Sections 3.1.1.3, 3.2.1.3, 3.3.1.3, and 4.1.1.3 contain the labeling requirements, which are summarized under various headings on pages 99-102 and page 105. Sections 3.2.3.1.3, 3.3.3.1.3, and 4.1.3.1.3 contain the covert channel requirements, which are summarized on page 97.

INTERP94 National Computer Security Center, *The Interpreted TCSEC Requirements*, (quarterly).

The following Interpretations are relevant to MAC:

I-0239	Subject access revocation after change in user clearance
C1-CI-02-86	Server

The following Interpretations are relevant to labeling:

I-0003	Access validation after object label change
I-0007	Assigning device level range
I-0022	One set of banner pages around multiple outputs
I-0039	Multilevel printers and page labeling
I-0040	Requirements for overwrite label capability
I-0253	Default page marking format
I-0275	Single-level printers and page labeling
C1-CI-05-84	Exportation to Multilevel Devices
C1-CI-01-85	Device Labels
C1-CI-01-88	Exportation of Labels
C1-CI-03-89	DAC Public Objects

The following Interpretations are relevant to covert channels:

C1-CI-02-84	Security Testing
C1-CI-07-84	Audit

Gasser88 Gasser, M., *Building a Secure Computer System*, Van Nostrand Reinhold Co., N.Y., 1988.

Chapters 6.3 and 6.4 talk about MAC. 6.5 talks about integrity and protection mechanisms. 7.2 talks about covert channels and 11.3 talks about protected subsystems.

Blotcky86 Blotcky, S., Lynch, K., and Lipner, S., "SE/VMS: Implementing Mandatory Security in VAX/VMS," *Proceedings of the 9th*

Module Eight

sensitivity label? (b) What part of the sensitivity label is output?
(c) Where is this output posted?

22. (a) How does the TCB designate the minimum and maximum sensitivity levels of a device? (b) List the ways these designations can be changed. (c) List the users who can invoke these mechanisms.
- 23 List the circumstances under which the TCB allows input or output of data that falls outside a device sensitivity range.

2.10 MANDATORY ACCESS CONTROL

B1:

1. Define the MAC policy for various access modes such as read, write, append, delete.
2. (a) Does the system use the sensitivity labels to enforce the MAC? (b) If not, what information is used to make the MAC decisions?
3. (a) List the subjects, objects, and circumstances under which the MAC policy is not enforced. (b) Why is it not enforced in these cases?
4. In what sequence does the system perform access mediation? (An example sequence might be a. privileges that supersede MAC and DAC, then b. check for DAC, then c. check for MAC.)
5. (a) Does the TCB support system-low and system-high sensitivity levels? If yes, how can they be (b) designated and (c) changed? Who can invoke the functions to (d) designate and (e) change them? How are these levels used by the system in (f) various labeling functions and (g) MAC decisions?

2.12 MODELING AND ANALYSIS

B2:

10. What tools, techniques and methodologies are used to identify, analyze, calculate, and reduce the bandwidths of covert channels?

2.13 OTHER ASSURANCES

C1:

5. (a) List separately the functions that can be performed by each of the trusted users (e.g., operator, security administrator, accounts administrator, auditor, systems programmer). (b) For each of these persons/roles, list the system data bases that can be accessed and their access modes. (c) Also list the privileges provided to each of these roles.
6. (a) How does the TCB recognize that a user has assumed one of the above- mentioned trusted roles? (b) Which of the above- mentioned functions can be performed without the TCB recognizing this role?

Module Eight

or (d) the actual labeling? (e) If so, what is the role of the person involved (e.g., system administration, system operator)? (f) Does this labeling require special privileges? (g) If so, what are those privileges?

8. (a) Who can change the labels on a subject? (b) How?
9. (a) Who can change the labels on an object? (b) How?
10. How are the labels associated with objects communicated outside the TCB?
11. (a) How does the system designate each device to be single-level or multilevel? (b) List the ways this designation can be changed. (c) List the users who can perform this designation.
12. (a) How does the TCB designate the sensitivity level of a single-level device? (b) List the ways this designation can be changed. (c) List the users who can do this.
13. (a) How does the TCB export the sensitivity label associated with an object being exported over a multilevel device? (b) What is the format for the exported label? (c) How does the TCB ensure that the sensitivity label is properly associated with the object?
14. (a) What mechanisms are available to specify the human-readable print label associated with a sensitivity label? (b) Who can invoke these mechanisms?
15. (a) Is the beginning and end of each hardcopy output marked with the human-readable print label representing the sensitivity level of the output (i.e., does each hardcopy output have the banner pages)? (b) What happens if a banner page output is longer and/or wider than a physical page?
16. (a) Is the top and bottom of each hardcopy output page marked with the human-readable print label representing the sensitivity level of the output? (b) What happens if the print label is wider and/or longer than the space available for the top and/or the bottom
17. How does the TCB mark the top and bottom page of non-textual type of output such as graphics, maps, and images?
18. (a) How can these markings be overridden? (b) Who can override the markings?
19. How can an operator distinguish the TCB generated banner pages from the user output?

B2:

20. (a) How does the TCB notify or acknowledge a change in the sensitivity level associated with an interactive user? (b) Is the user notification posted on the user terminal? (c) How immediate is this change?
21. (a) How does a user query the system TCB for his or her current

Module Eight

11. (a) List other shared resources which are not protected by the MAC mechanism. (b) Why are they not protected? (c) Describe the mechanisms that are used to isolate and protect these resources.

2.4 SOFTWARE

The TCB software consists of the elements that are involved in enforcing the system security policy. Examples of TCB elements include: kernel, interrupt handlers, process manager, I/O handlers, I/O manager, user/process interface, hardware and command languages/interfaces (for system generation, operator, administrator, users, etc.). The security kernel is the hardware, firmware and software elements of the TCB that are involved in implementing the reference monitor concept, i.e., the ones that mediate all access to objects by subjects.

C1:

6. List all the privileges a process can have. Include the privileges based on the process or user profile, process or user name, or process or user identification.
7. How are a process's privileges determined?

B1:

21. How is a process' sensitivity level determined?

2.9 LABELS

B1:

1. (a) How many hierarchical sensitivity classifications (such as unclassified, confidential, secret, top-secret), does your system provide for? (b) What mechanisms are available to define the internal/storage and external/print format? (c) What mechanisms are available to change them? (d) Who can invoke these mechanisms?
2. (a) How many non-hierarchical sensitivity categories (such as FOUO) does your system provide for? (b) What mechanisms are available to define the internal/storage and external/print format? (c) What mechanisms are available to change them? (d) Who can invoke these mechanisms?
3. (a) What is the internal TCB storage format of the sensitivity label? (b) If different for different subjects or objects, give all formats.
4. For each type of subject, where is the subject sensitivity label stored?
5. For each type of object, where is the object sensitivity label stored?
6. (a) List the subjects and objects that are not labeled. (b) Why are they not labeled? (c) How are these subjects and objects (c) accessed and (d) controlled?
7. (a) How is imported data, labeled? (b) How is this label determined? Is a human being involved in (c) the determination

Module Eight

the object. As can be seen, all MAC decisions are controlled by the TCB and cannot be affected by a normal user.

Relevant Trusted Product Evaluation Questionnaire Questions

2.1 SUBJECTS

A subject is an active entity in the system, generally in the form of a person, process, or device that causes information to flow among objects or changes the system state. A subject can be viewed as a process/domain pair whose access controls are checked prior to granting the access to objects.

C1:

4. (a) What are the security attributes of a subject? (Examples of security attributes are user name, group id, sensitivity level etc.) For each type of subject in your system (i.e., user, process, device, etc.), what mechanisms are available to (b) define and (c) modify these attributes? (d) Who can invoke these mechanisms?
5. (a) What are other privileges a subject can have? (Examples of such privileges are: super user, system operator, system administrator, etc. Your operating system may assign numerous other privileges to the subjects, such as the ability to use certain devices.) For each type of subject in your system, what mechanisms are available to (b) define and (c) modify these privileges? (d) Who can invoke these mechanisms? (e) Provide a list of subjects within the TCB boundary and (f) the list of privileges for each of them.
6. When a subject is created, where do its (a) security attributes and (b) privileges originate, i.e., how are the security attributes and privileges inherited?
7. List the subjects, if any, which are not controlled by the TCB.

2.2 OBJECTS

An object is a passive entity that contains or receives information. Access to an object potentially implies access to the information it contains. Examples of objects are: records, blocks, pages, segments, files, directories, directory trees, and programs, as well as bits, bytes, words, fields, processors, video displays, keyboards, clocks, printers, network nodes.

C1:

1. Provide a list of objects within the TCB (e.g., authentication database, print queues).

B1:

9. List the objects in your system that are protected by the Mandatory Access Control (MAC) mechanisms.
10. (a) List the objects that are not protected by the MAC mechanism. (b) Why are they not protected? (c) Describe other mechanisms used to isolate and protect objects.

Module Eight

names and a low-level process could read these names, thus allowing an indirect write down from the high to the low process. This channel could be eliminated by disallowing the low-level process from reading the filenames, or even knowing that the high-level files exist. There are many types of covert channels, most not as obvious as the example, and more difficult to eliminate.

The covert channel requirements in the TCSEC do not begin until B2 even though covert channels become an issue as soon as MAC is introduced at B1. The decision on whether a violation of policy is a flaw or covert channel is one that must be negotiated by the vendor and team involved. Design flaws cannot just be called a 'covert channel' and ignored because of the lack of a requirement and the refusal to fix a flaw.

The first explicit covert channel requirement occurs at B2. At B2, a thorough search for covert storage channels must be made. It is not required that all covert channels found must be eliminated. For each channel found, the bandwidth of the channel must be determined and the steps that can be taken to reduce the bandwidth or eliminate the channel identified. At a minimum, covert channels must be documented. Beyond certain bandwidths, as defined in the covert channels guidance in the TCSEC, they must also be audited. This guidance also defines maximum thresholds that the covert storage channels must not exceed. At B3, the requirement is expanded to include a thorough search for covert timing channels as well. Finally, at A1, formal methods are required to be used during the covert channel analysis process.

The following is the current guidance supported by the NSA Trusted Product Evaluation Program regarding covert channel bandwidth.

“In an evaluated system, covert storage channels that are:

greater than 100 bits per second shall not exist;

10 to 100 bits per second shall be documented, and shall have their real or potential uses audited;

1 to 9 bits per second shall be documented and may be audited;

less than 1 bit per second may be documented and may be audited.”

Summary

Users are typically assigned a maximum sensitivity label (in some cases a range) by the system security officer, which defines the limits at which that user can operate. The maximum sensitivity label should, of course, have a hierarchical sensitivity level no greater than the user's clearance and its category set should include only categories for which the user is authorized access. When a user logs into the system, the user indicates a desired sensitivity label for the current session. This sensitivity label must be dominated by the clearance (or be within the range) assigned to that user and within the range allowed for the terminal. Then, whenever the user's process performs a read or write, it is constrained by the MAC read and write controls previously described. Whenever the process creates an object, the TCB must ensure that the object has a sensitivity label that dominates the user's current sensitivity label. Frequently, the TCB simply assigns the user's current label to

Module Eight

As mentioned earlier, subjects (usually processes) must have a sensitivity label. This sensitivity label defines the process' hierarchical level (e.g., Unclassified, Secret, Top Secret, etc.) and its category set (e.g., NATO, No Foreign, etc.). All objects (files, directories, etc.) are also required to have a sensitivity label. The access modes defined (read, write) are constrained as described above. The sensitivity labels are partially ordered by the dominance relationship defined by: "sensitivity label S1 is said to dominate sensitivity label S2 if the hierarchical classification of S1 is greater than or equal to that of S2 and the non-hierarchical categories of S1 include all those of S2 as a subset" [TCSEC85]. Using this term, the TCSEC requirement above states that a subject can only read objects that it dominates and can only write to objects that it is dominated by. These constraints are reflected in the Bell & LaPadula model's simple security property and *-property (refer to Module 5).

For some objects that exhibit the characteristics of a queue or stack, the act of reading the object causes the state of the object to change (e.g., popping an entry from a stack deletes the entry). In cases such as these, the evaluation community insists that for the read operation to be allowed, the subject would need not only read permission (i.e., subject dominates object) but also write permission (e.g., object dominates subject). Therefore, for objects that have destructive reads, the operation can be permitted only if the subject sensitivity label equals the object sensitivity label.

The main difference between MAC and DAC is that users cannot modify the mandatory access attributes (labels) of objects, while users can modify the discretionary access attributes of objects.

Covert Channels

The discussion above centered around MAC on explicit reads and writes. Information can also be compromised in violation of the MAC policy by covert channels. A covert channel as defined in the TCSEC is:

"A communications channel that allows a process to transfer information in a manner that violates the system's security policy."

There are two types of covert channels. They are covert storage channels and covert timing channels.

Covert storage channel: A covert channel that involves the direct or indirect writing of a storage location by one process and the direct or indirect reading of the storage location by another process.

Covert timing channel: A covert channel in which one process signals information to another by modulating its own use of system resources in such a way that this manipulation affects the real response time observed by the second process.

Covert channels can be used to circumvent the MAC policy by allowing a high level process to signal a low level process, effectively causing a write down. This is a threat that could be used by a Trojan horse to leak information out of a classified system, even though a MAC policy is in force. An example of a covert channel would be the use of file names. If filenames are not protected at the level of the file itself, a high-level process could create high-level files of various

Module Eight

with the same sensitivity label can be stored, displayed, printed, etc., on that device. Multilevel devices can support a range of sensitivity labels. At B2 and above, devices are given a minimum sensitivity label and a maximum sensitivity label, constraining the range over which a multilevel device may operate, or defining the range over which a single-level device can operate a single level at a time. Multilevel devices can be used to store, display, print, etc., any data that dominates the minimum sensitivity label and is dominated by the maximum sensitivity label. Single-level devices, like terminals, must operate at a single sensitivity label at any one time unless specifically designed and trusted to perform multilevel functions (in which case they are really multilevel devices). These devices have both an assigned range and a current sensitivity label. The current sensitivity label can be changed between user sessions.

MAC

Mandatory access control in the paper world restricts reading to those users that have the appropriate clearance. It does not restrict a person from writing information at a lower classification level after reading a document of a higher classification level. This is because the person is trusted to not divulge any higher level information when writing a document with a lower classification level. For example, the President of the United States is still permitted to make public speeches even though he knows a large number of highly classified state secrets. However, within a computer system, the paper world's rules are incomplete. Within a computer system, users execute a great deal of software of which they may have little knowledge. This software cannot, and should not, be trusted by the user or the TCB. When running untrusted software at a certain security level, the software cannot be trusted to not divulge information (i.e., write the information to an object with a lower sensitivity label). Thus, along with ensuring that users don't read objects above their clearance, the MAC mechanism must enforce a no write down policy to ensure that untrusted software does not compromise any information.

The TCSEC does not impose a MAC requirement for systems at the C1 or C2 assurance level. The first MAC requirements begin at B1 where subject and object labels are required and all accesses between subjects and objects must be mediated according to the MAC policy. As mentioned earlier, a relationship between sensitivity labels is enforced as part of MAC. This relationship, known as dominance, is defined in the following quotation from the TCSEC MAC requirement.

"A subject can read an object only if the hierarchical classification in the subject's security level is greater than or equal to the hierarchical classification in the object's security level and the non-hierarchical categories in the subject's security level include all the non-hierarchical categories in the object's security level. A subject can write an object only if the hierarchical classification in the subject's security level is less than or equal to the hierarchical classification in the object's security level and all the non-hierarchical categories in the subject's security level are included in the non-hierarchical categories in the object's security level."

Module Eight

At B1, the TCSEC requires that a sensitivity label be assigned to every storage object. For classes B2, B3 and A1, a sensitivity label must be assigned to every resource. Simply labeling the objects in a system is not enough to allow a TCB to implement a MAC mechanism. As described by the TCSEC MAC requirement, a comparison is performed between the subject's label and the object's label. Therefore, all subjects in the system must also have a sensitivity label. Beginning at class B2, the TCSEC "Subject Sensitivity Label" requirement also specifies that a user be able to ask the TCB for the value of that label.

Label Integrity

The TCSEC "Label Integrity" requirement explicitly states that the sensitivity labels assigned to subjects and objects must be accurate. That is, only the TCB may be permitted to set their values. It also requires that labels exported from the TCB be unambiguously associated with an object. While both of these concepts are rather obvious, they are the backbone of the entire MAC mechanism. Without accurately labeled subjects and objects, any MAC decision would be meaningless.

Import / Export of Labels

Objects which are imported or exported must also be protected. The TCSEC defines these import/export requirements. There are many issues that have to be dealt with in this area: compatibility requirements between labels from different systems, using the same or another vendor's operating system, label integrity while the label is outside the control of a TCB, import of unlabeled information, etc. The most basic point of these requirements, however, is that the association of a label with an exported or imported object must be handled in such a way that the object is always labeled accurately. Refer to the Exportation of Labeled Information, Exportation to Multilevel Devices, and Exportation to Single-Level Devices requirements in the TCSEC.

Human-Readable Labels

There are also requirements for labeling human-readable output, such as printed output, maps, plots, etc. For each of these output devices, a format must be defined that allows the security label of the output object to be prominently displayed on the front of the output media when an object is output in this form.

Another part of this requirement explicitly states that the "override of [human-readable] markings shall be auditable by the TCB." The evaluation community has interpreted "override" as being equal to "overwrite" when it comes to human-readable labels. That is, if the region of a printout in which the TCB top and bottom page label would be printed is writable by a user, the user has the ability to overwrite (and thus override) the printing of page labels by the TCB.

Device Labels

The TCSEC also requires that devices be labeled starting at B2. Devices can be single-level or multilevel. On single-level devices, only information labeled

Module Eight

Mandatory Access Control and Labels

This module describes the concept of mandatory access control (MAC). In this context it introduces sensitivity labels and their use by the MAC mechanism. The TCSEC requirements for MAC and labels are given and then some actual implementations of MAC mechanisms are described, as well as some of the implementation requirements created by the need for MAC.

Module Learning Objectives

The material in this module can be read independently, but it draws somewhat on the material presented in Modules 5 and 6. Upon completion of this module, the student should:

1. Understand MAC.
2. Understand the implications of the requirement for MAC, including the need for labels and the problem of covert channels.
3. Understand the TCSEC MAC requirements.
4. Understand the TCSEC Label requirements.
5. Understand the TCSEC Covert Channel requirements.
6. Be familiar with some implementations of MAC.
7. Be familiar with issues of the implementation that require special attention.

Overview

The concept of MAC is not as well understood by many members of the community as the familiar mechanisms of discretionary access control. MAC is, as the name implies, an access control mechanism enforced by the TCB that must mediate all accesses between subjects and objects in a defined manner. It is mandatory in the sense that the assignment and use of MAC information is not under the influence of a user's discretion. The user does not have any discretionary ability to modify any mandatory access rights associated with subjects or objects. The TCB makes a MAC decision based upon a subject's sensitivity (i.e., clearance) and an object's sensitivity (i.e., classification). This module provides a very brief description of the TCSEC's requirements related to labeling, MAC, and covert channels.

Labeling

Sensitivity labels are central in the enforcement of MAC. The make-up of a sensitivity label is described in the TCSEC MAC requirement. It says that a sensitivity label has two parts, a hierarchical level and a set of non-hierarchical categories. The MAC requirement also describes the way in which labels are compared. The method of comparison defines a relationship between labels referred to as dominance. Dominance is described later in this module. The TCSEC labeling requirements address what must be labeled. They cover subject and object labeling, label integrity, exportation of labeled information, device labeling, and labeling of human readable output. Each of these topics is briefly described in this module and in the required readings.